

UNITED STATES PATENT APPLICATION

FOR

APPROACHES FOR MIGRATING PORTAL OBJECTS
FROM A SOURCE INSTALLATION TO A TARGET INSTALLATION

INVENTORS:

RAVISHANKAR RAMANATHAN
ARUN SHIVASWAMY
BALAKRISHNAN JAGDEESAN

PREPARED BY:

HICKMAN PALERMO TRUONG & BECKER LLP
1600 WILLOW STREET
SAN JOSE, CALIFORNIA 95125
(408) 414-1080

ASSIGNEE:

ORACLE INTERNATIONAL CORPORATION
500 ORACLE PARKWAY
REDWOOD SHORES, CA 94065

"Express Mail" mailing label number EV322192892US
Date of Deposit February 26, 2004

APPROACHES FOR MIGRATING PORTAL OBJECTS
FROM A SOURCE INSTALLATION TO A TARGET INSTALLATION

CLAIM OF PRIORITY

[0001] This patent application claims priority from U.S. Provisional Application Serial No. 60/482,910, entitled “EXTENSIBLE FRAMEWORK IN AN APPLICATION DESIGN TOOL TO MIGRATE CUSTOMIZABLE AND USER-BUILT “APPLICATION COMPONENTS” AND “DESIGN ELEMENTS” FROM ONE INSTALLATION TO ANOTHER,” filed June 25, 2003, by Ravishankar Ramanathan, Arun Shivaswamy, and Balakrishnan Jagdeesan, the contents of which are incorporated herein by reference for all purposes.

FIELD OF INVENTION

[0002] The invention generally relates to computer systems, and, more specifically, relates to approaches to migrate portal objects from a source installation to a target installation.

BACKGROUND OF THE INVENTION

[0003] The approaches described in this section are approaches that could be pursued, but not necessarily approaches that have been previously conceived or pursued. Therefore, unless otherwise indicated, the approaches described in this section are not prior art to the claims in this application and are not admitted to be prior art, merely by inclusion in this section.

[0004] The creation of software applications is becoming more and more complex. For example, hundreds of software engineers spanning the globe may simultaneously

work on the production of a single software application, or may simultaneously work on various related software applications. Additionally, the work on the software application may continue throughout the day and night. For example, engineers in California may work on a certain portion of a software application during a day. At the end of the day, the application may then be e-mailed to another team of engineers working in another country, such as India or Japan, so that another other team of software engineers may continue to work on the application. In this manner, the creation of the software product may continue around the clock.

[0005] Additionally, software applications are often created incrementally, and must be moved from one environment to another environment. For example, a software application may need to be moved from a production environment to a testing environment. However, when engineers are scattered across the globe, and production continues around the clock, various challenges arise in maintaining integrity and coalescence among single and multiple applications. For example, metadata may be changed by an engineer working on an application in California, but, if an engineer in India or another country is working on the same application or a related application, metadata in a portal object that is sent from the engineer in California to the engineer abroad may be corrupted or inconsistent.

[0006] During software production processes, various installations (e.g., portals) may be involved. In each of the installations, portal objects are created. Portal objects may include application components and/or design elements. The application components used by an application may include, for example, Forms, Reports, Charts, etc. The design elements used by an application may include, for example, Pages (Deciding on the

Layout and Regions), Styles, Templates, etc. Portal objects may be migrated for various purposes, such as, for example, moving a portal object to another installation so that another software engineer may alter or test the portal object in an incremental software development and testing environment.

[0007] When a portal object is created, metadata concerning the portal object is often stored in a portal table. An installation may have as many as 500 or more tables that comprise the architecture of the software application. The portal tables are used to hold metadata of different kinds of objects. There is often a need to migrate such portal objects from one installation of portal to another installation of portal. For example, if 25 distinct portal objects, each of a different type, need to be migrated to a different portal instance, then a subset of the more than 500 tables, which a user believes may encapsulate the metadata of all twenty-five of the different objects, may be imported or exported. However, complications arise because during the migration, the metadata of objects that already exist on the target portal may be corrupted during the import.

[0008] One prior approach to maintaining consistency across applications is versioning. Different versions of the application, including a master version of the application, were maintained. When an engineer completed his or her version of the application, it was copied into the master version. This approach is cumbersome and does not always result in consistent data across applications.

[0009] Most of the previously existing approaches for migrating portal objects employ hard-coded and ad-hoc logic to copy metadata from a source portal to target portal. The fact that existing solutions must be hard-coded and prepared in an ad-hoc manner does not allow parallel developments to occur in an efficient fashion. Moreover,

as mentioned above, in existing approaches for migrating portal objects, metadata is often corrupted. Hence, there is a need for a framework to be designed in such tools to facilitate the migration of portal objects in a systematic way from a source installation to a target installation.

BRIEF DESCRIPTION OF THE FIGURES

[0010] The disclosed subject matter is depicted by way of representative examples, and not by way of limitation, in the illustrations of the accompanying figures. The figures include:

[0011] FIG. 1 is a flowchart depicting an embodiment of a method for migrating portal objects from a source installation to a target installation.

[0012] FIG. 2 is a flowchart depicting object selection and metadata extraction from portal system tables into the corresponding transport tables.

[0013] FIG. 3 is a block diagram depicting a system upon which embodiments of the disclosed approaches may be implemented.

[0014] FIG. 4 is a block diagram depicting portal objects to be transferred from a portal system table to a transport table.

DETAILED DESCRIPTION OF THE INVENTION

[0015] In the following description, for the purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be apparent, however, that the present invention may be practiced without these specific details. In other instances, well-known structures and devices are shown in block diagram form in order to avoid unnecessarily obscuring the present invention.

[0016] Approaches are disclosed for migrating objects, created for an application by an application design tool, from a source installation to a target installation. The approaches may be used, for example, to migrating any type of portal object from a source installation to a target installation. The portal objects selected to be migrated from a source installation to a target installation may be referred to as the “migrating objects”. In an embodiment, a logical container called a “transport set” stores the migrating objects. The metadata of the migrating objects may be stored in infrastructure tables. As shall be described in greater detail hereafter, during the migration, the data in the transport tables (i.e. the application specific data) and the metadata in the infrastructure tables (about the migrating objects) are exported and/or imported to the target installation. In an embodiment, scripts are generated to perform export and/or import functions.

[0017] A number of benefits may be realized using the migration techniques described herein. For example, in one embodiment, an infrastructure is provided that facilitates the migration of portal objects among different installations in a manner that provides improved consistency across installations and reduces corruption. Techniques

are also provided to migrate portal database providers, page groups, Pages / Portlets, etc. from one installation to another installation. Using the techniques described herein, one may migrate provider registration information for web and database providers whose portlets appear on a page or page group that is being migrated. The techniques allow for the explicit migration of shared objects or shared components.

[0018] A technique is provided to selectively mark objects for export and save them to a transport set. The disclosed approaches are also useful in the creation of backups of portal objects or installations.

[0019] The foregoing benefits, which derive from certain embodiments, provide an efficient framework for the incremental development of large applications in complex software development environments, including, in particular, staging-to-production development environments that span the globe and feature around-the-clock development by different teams of software engineers.

FUNCTIONAL OVERVIEW

[0020] A framework is described herein for facilitating the migration of objects of an application, created using an application design tool, from one installation to another. Typically, as the application design tool is used to create objects for an application, the application design tool will store metadata about the objects in system tables, which typically reside within a database system. The migration framework includes a mechanism for maintaining a second set of tables, called transport tables, for holding the relevant metadata of the objects that are to be exported to another installation. According to one embodiment, the transport tables are mirror images of the corresponding dictionary tables of the application design tool.

[0021] As shall be described in greater detail hereafter, the framework also includes a logical container called a “Transport Set” to hold the specific objects to be migrated. The framework also provides mechanism, such as a user interface, through which various “Components” that need to be migrated can be identified and added to the transport set. Also, the framework allows different modes of migration, such as “Overwrite”, “Reuse”, and “Check”. Through specifying a mode, the end-users can determine the behavior in the target installation when such components are migrated, and there is a parallel development on the other side.

[0022] Approaches are disclosed for migrating portal objects from a source installation to a target installation. In one embodiment, the approaches for migrating portal objects include approaches for migrating customizable and user-built application components and design elements from one installation to another.

[0023] According to one embodiment, the migration process involves the creation and maintenance of a set of tables, called transport tables, for holding the metadata to be migrated from one installation to another. Metadata from source installation tables, such as dictionary or system tables, may be stored in the corresponding transport tables (with some preferable extra columns) to facilitate migration to the target installation.

[0024] In an embodiment, a logical container called a “transport set” stores the specific portal objects to be migrated. The transport set provides a way through which various portal objects that need to be migrated may be identified. The transport set is stored in infrastructure tables. According to one embodiment, that transport tables are actually mirror tables of the application specific system tables.

[0025] By using a transport set, a number of generic notions and modes, such as “Overwrite”, “Reuse”, and “Check” may be implemented. Using these notions and modes, users may specify allowable or desired data protection modes (e.g., read/write privileges) in the target installation when the migrating objects are migrated. The specified data protection modes dictate how (1) the migrated data and (2) the pre-existing data in the target installation, are reconciled. Based on the data protection modes, data consistency across installations can be maintained.

[0026] Several of the migration techniques described herein employ: i. data extraction; ii. script generation; and iii. data merge. The data merge process, in turn, may involve a resolution operation. Each of these steps shall be described in greater detail hereafter.

[0027] According to one embodiment, migration begins with a step of selecting portal objects for migration. As mentioned above, each portal object that is selected for transport (each “migrating object”) is put into a container object called a transport set.

[0028] After the transport set is finalized, the transport set is subject to a data extraction step. In the data extraction step, metadata of the migrating objects is pulled from the portal system tables or dictionary tables and stored in transport tables.

[0029] In an embodiment, at the end of the data extraction process, a script is generated for use in the export and/or import of the migrating objects. The script is made available for download. A purpose of this script is to perform validation on the transport set and invoke the export and/or import utilities internally, based on the mode chosen for running the script. The script may be run to generate a dump file for export. In an

embodiment, a user may input parameters into the script. The dump file contains, in essence, a dump of the transport set containing portal objects selected for migration.

[0030] After the dump file is created, the same script may be run in import mode by specifying the dump file as the input, along with any other parameters used by the script. Execution of the script in import mode migrates the transport set to the target installation.

[0031] Even after the import of the migrating objects into the target installation, the migrating objects may still not be ready for use in the target installation, as the migrating objects may not yet be integrated into the system tables or dictionary tables of the target installation.

[0032] In another step, the migration may involve a merge operation. In the merge operation, the data residing in the transport tables is transferred to the system or dictionary tables of the target installation. This step may include processes for resolving anomalies in data that resides in the transport set, and/or for improving or ensuring consistency among transport set data and data that is pre-existing in the target installation.

[0033] Rather than previous ad-hoc, hard-coded, inconsistent approaches to migrate portal objects among installations, the disclosed approaches provide a systematic and extensible framework to migrate portal objects. In addition to migration, this framework facilitates the backup and recovery mechanisms for various portal objects that may be built by various users at different installations.

[0034] The disclosed approaches also allow for parallel development to be performed, improved consistency, and reduced data corruption in across multiple installations, such as the various installations involved in a migration from a staging environment to a production environment.

PORTAL IMPLEMENTATION OVERVIEW

[0035] The migration framework described herein may be employed in a portal environment. According to one embodiment, when employed in a portal environment, the framework includes a portal export/import utility maintains transport tables at each portal installation. According to one embodiment, every transport table is an approximate a mirror image of its corresponding portal system table. Every portal object which has been identified for export is put into a container type of object called a transport set.

[0036] After the transport set is finalized, the transport set is submitted for the data extraction process. In this step, all the relevant metadata of the portal objects comprising the transport set are pulled from the portal system tables to the respective transport tables. As such, the portal system tables contain metadata of all the objects of portal, and the framework selectively picks metadata of those objects placed in the transport set only.

[0037] At the end of the data extraction process, a script is made available for download. In one embodiment, execution of the script performs validations on the transport set and invokes the import and/or export utilities of the database server that manages the tables in question. The import/export utility is invoked internally based on the mode (import or export) chosen for the running the script. To extract the data from the source installation, the script is executed in EXPORT mode to generate a dump of the transport set containing those selected objects chosen for export.

[0038] Once the dump has been created, the same script can be executed in IMPORT mode by specifying the dump as the input. Execution of the script in import mode, brings in the transport set to the target portal instance.

[0039] To make the portal objects useable at the target portal instance, a merge operation is performed to make the metadata for the transport set an integral part of the target portal's system tables. During the merge operation, the data residing in the transport tables are pushed into system tables of the target portal. After this is done successfully, the migrated objects are available to be viewed and accessed at the target portal installation.

MIGRATION TECHNIQUES AND TRANSPORT SETS

[0040] FIG. 1 is a block diagram depicting an embodiment of system 100 upon which the migration techniques described herein may be implemented. System 100 includes a source dictionary table 120. Source dictionary table 120 holds the metadata for portal objects in a source installation. A single dictionary table 120 is shown for the purposes of simplicity. However, at the source installation, the metadata for portal objects may be spread across any number of dictionary and/or system tables.

[0041] In the illustrated embodiment, system 100 also includes a source transport table 140. In one embodiment, source transport table 140 is a mirror of dictionary table 120, with zero or more additional columns. The metadata of portal objects that are selected for migration to a target installation is extracted from source dictionary table 120 and exported to source transport table 140.

[0042] In an embodiment, the metadata stored in source transport table 140 is then migrated to the target installation by importing the metadata from source transport table 140 into target transport table 160. In an embodiment, the transfer of data from source transport table 140 to target transport table 160 is performed through the use of a dump

file. Specifically, the import operation is facilitated by generating and running the script to create a dump of source transport table 140. Once the dump is created, the data is imported from the dump into the target transport table 160.

[0043] Metadata for the migrating objects may then be imported to target dictionary table 180. A data merge routine may be used to merge the metadata in source transport table 160 with any pre-existing data in target dictionary table 180. Additionally, prior to the data merge, the metadata in target transport table 160 may be examined for anomalies. If anomalies exist, measures may be applied to resolve the anomalies prior to the data merge.

[0044] As discussed above, in an embodiment, only one transport set is needed, and can be used to facilitate the migration of any number of migrating portal objects.

[0045] FIG. 2 is a block diagram of system 200 depicting portal objects to be transferred from source table 210 in a source installation to a target installation. In the illustrated embodiment, source table 210 is a portal system table, a dictionary table or any other table in a source installation used to store metadata for portal objects in the source installation. For the purpose of explanation, an embodiment is described in which the metadata for a target object resides in a single system table. However, an object may have its metadata stored in more than one system table.

[0046] Portal objects may be selected for migration, as described above. In the example shown in system 200, any of portal objects 215 – 265 may be selected by a user for migration. In the embodiment depicted, the user selects portal objects 225 and 255 for migration to a target installation. Accordingly, portal objects 225' and 255' are stored in transport set 260.

[0047] Once the selected migrating objects are stored in transport set 260, portal objects 225' and 255' may be imported into a target installation and merged with any pre-existing data in the target installation.

[0048] In an embodiment, implicitly referenced portal objects may also be selected. For example, suppose that a user selects only portal object 225 for export. If, in order to maintain consistency due to codependences, it is also required implicitly that portal object 255 must be exported (due to foreign-key / master-detail / parent-child relationships), then portal object 255 may be implicitly selected for migration.

EXEMPLARY APPROACHES

[0049] According to one embodiment, the migration of portal objects involves several phases, including: i. data extraction; ii. script generation; and iii. data merge. These phases are described in greater detail below.

[0050] (i) Data Extraction: Metadata to be migrated is extracted from a selected portal object. In an embodiment, each tool component handles the extraction of data from their dictionary tables and storage in the transport tables. Thus, extensible, the various tool components can extract using their own specific logic/dictionary table structures.

[0051] (ii) Script Generation: In script generation, a script is generated for the user to enable extracting this metadata to the OS files. Underneath, the script can make use of RDBMS utilities such as export and import. The user runs the script to store everything onto OS files, and to ultimately move the metadata into the transport tables of the target installation.

[0052] (iii) Data Merge: In an embodiment, individual tool components register their own logic for merging the migrated metadata to the target dictionary tables, as appropriate. In this phase, any predefined modes for each and every component being merged could be overridden by users. Prior to the data merge, the metadata in a target transport table may be examined for anomalies. If anomalies exist, measures may be applied to resolve the anomalies prior to the data merge.

OBJECT SELECTION

[0053] FIG. 4 is a flowchart of method 400 for migrating one or more objects from a source installation to a target installation. Step 410 of method 400 involves selecting one or more portal objects to migrate from a source installation to a target installation.

[0054] The portal objects may be selected for migration using a graphical user interface displayed on a monitor. In this regard, from a portal navigator or bulk actions screen, a user may click on an “Export” link that is next to a visual representation of a portal object. Clicking the export link that is next to a visual representation of a portal object will select the corresponding portal object for migration. Clicking on the export link may open a transport set user interface to allow a user to create a new transport set or to add the selected portal object(s) to an existing transport set.

[0055] The type of objects that may be exported by a user depends on the class of the user. In this regard, different categories of export links will be available to users based on the user’s class (e.g., regular user, administrator, etc.). Export links may correspond to specified page groups, page group objects, portal database providers, portal database provider components, or other components. Additionally, if a user has super

administrator privileges, the user may click on an export link corresponding to shared portal objects, in an embodiment.

[0056] If a user selects the “create new transport set option”, the user is presented with other options, such as naming the transport set or choose a default name for the transport set. Next, the user is taken to a manifest screen that displays the manifest for the portal object(s) selected for migration.

[0057] In an embodiment, the manifest screen displays an option whereby a user may select the export of object access control lists (ACL). The user may also optionally select preferences for users and groups associated with the ACLs. In addition to the fact that the user may explicitly select portal object(s) for export, other portal objects that the explicit selection references (such as styles, categories, etc.), and external references (such as web and database providers), other page groups and portal database providers, etc., may also be selected. As a default, these options may initially be deselected (e.g., no ACLs or preferences associated with the ACL’s will be migrated).

[0058] Each representation of an explicit portal object in the transport may be associated with a hyperlink that displays the manifest for the explicit portal object. Therefore, if two portal objects are added to the transport set, the references for each portal object may be independently identified using the hyperlink for that portal object.

[0059] The transport set could then be saved for export at a later time, or the transport set may be processed immediately.

[0060] As stated above, step 410 optionally comprises adding an object to an existing transport set. If this is the case, then the user selects to add an object to an existing

transport set, rather than creating a new transport set and adding the object to the new transport set.

[0061] The creation of the transport set (e.g., the transport tables) and the addition of objects to the transport set is facilitated by functions in an application layer, which in turn perform all the database related operations through a database layer.

Tables Used in Migration

[0062] Various tables are used in implementing migration techniques that use transport sets. According to one embodiment, a master table stores: metadata about a transport set, data that indicates whether to export object ACLs, data that indicates preferences for users/groups in ACLs, and data that indicates default options for import including whether to ignore warnings. A detail table stores information about explicit and implicit portal object selections and external references. The detail table also stores information about the mode of import for explicit and implicit portal objects.

METADATA EXTRACTION

[0063] Method 400 also includes step 420. Step 420 includes extracting metadata about the selected portal objects from system tables into the corresponding transport tables in the source installation. In one embodiment, in step 420, extraction from the selected portal objects is initiated. Then, in step 330, the export-import infrastructure of the application initiates calls to sub-system export packages through standard routines that perform exports of data from selected portal objects.

[0064] In an embodiment, during step 420, a link is provided from the export-import user interface that allows the user to see the log of actions being performed. Additionally, the user interface allows a user to download specific scripts for transport functions that may be used in command lines to create a dump file for the transport set.

[0065] A number of tables are employed to implement steps 420 and 330. Specifically, as mentioned above, a master table stores: metadata regarding the transport set, data that indicates whether to export object ACLs, preferences for users/groups in ACLs, and default options for import, including whether to ignore warnings. A detail table stores information concerning explicit and implicit object selections and external references, mode of import for explicit and implicit object selects which are overwrite and reuse by default respectively, etc. A log table stores information about what happened to each migrating object, etc.

[0066] According to an embodiment, every system table that comprises the application being developed has a corresponding mirror table, called transport table. The transport tables may include some extra columns that are not included in the corresponding system table. These columns could typically store unique Ids to identify the transport sets and store the import mode of each migrating object.

SCRIPT GENERATION

[0067] Referring again to FIG. 4, at step 430 a script is generated which, when executed, will perform at least one of the aforementioned export steps and import steps. The user may invoke the script from command line, pass values to the parameters of the script, and create a dump file for the uniquely identified transport set.

[0068] The Windows NT and UNIX version of the script is generated by procedure(s) in a user interface package that embed a unique identifier for the transport set. The script invokes the Oracle Database's Export utility that exports the content of the transport tables (and the export-import utility tables) selectively. The script is common for both export and import, the parameters passed during invocation of the script and the values provided determine whether the operation is an export or import.

[0069] Various tables are involved in the implementation of step 430. For example, a master table stores metadata about transport set, whether to export object ACL's, preferences for users/groups in ACL's, default options for import including whether to ignore warnings. A detail table stores information about explicit and implicit object selections and external references, mode of import for explicit and implicit object selects which are overwrite and reuse by default respectively, etc. Various sub-system transport tables are also employed in step 430.

IMPORTING THE EXTRACTED METADATA

[0070] Method 400 includes step 440. Step 440 includes importing extracted metadata in the source transport table to a target transport table (i.e., to a transport table associated with the target installation.)

[0071] In an embodiment, if the metadata resides in a dump file created from the source transport tables, then the importing step may include invoking a script to move the contents of dump file(s) to the target transport tables.

[0072] Various packages are employed in the implementation of step 440. The Windows NT and UNIX version of the script is generated by procedure(s) in a

computer's user interface package and establishes a unique identifier for the transport set. This script is used during export to extract the contents from transport tables to dump file. The same script may used here to import the contents from dump file to transport tables.

[0073] In an embodiment, the script invokes the Oracle Database's Import utility that exports the content of the transport tables (and the export-import utility tables) selectively. As mentioned above, the script may be common for both export and import. The parameters passed during invocation of the script, and the values provided, determine whether the operation is an import or export.

[0074] A number of tables are employed in the implementation of step 440. For example, as mentioned above, the master table stores metadata regarding a transport set. The master table also indicates whether to export object ACLs, indicates preferences for users/groups in ACL's, and indicates default options for import including whether to ignore warnings.

[0075] A detail table stores information about explicit and implicit portal object selections and external references. The detail table also indicates modes of import for explicit and implicit portal object selections. Examples of modes of import include "overwrite" and "reuse", among others. Various subsystem transport tables are also employed.

RESOLVING ANAMOLIES

[0076] Method 400 also includes step 450. Step 450 involves resolving anomalies in metadata of the target transport tables and target system tables. Resolving anomalies may

include examining the metadata in objects and then correcting inconsistencies between the metadata in objects and metadata stored in files or other tables.

[0077] A call to start import for each sub-system will transfer control over to the sub-system which will then, resolve and merge data from transport tables to system tables, log information, and pass back control to export-import utility upon completion. A link may be provided from the export-import user interface that allows the user to see the log of actions being currently performed.

[0078] Step 450 may be initiated and performed at various points in method 400. In this regard, step 450 may be applied in respect of metadata in a transport set associated with a source installation and/or a transport set associated with a target installation. Step 450 may be applied at multiple instances in respect of each transport set, or it may be applied once in respect of a various transport sets (e.g., in respect of multiple transport sets in a single installation, or in respect of transport sets for different installations.)

[0079] In an embodiment, step 450 is performed prior to merging data from a transport table associated with a target installation into the system or dictionary tables of the target installation.

MERGING THE METADATA

[0080] Method 400 also includes step 460. Step 460 involves merging metadata from the target transport tables into system tables or other tables in the target installation. For example, step 460 may include merging metadata from a target transport table into a system table or dictionary table in the target installation.

[0081] In order to move the portal object data into tables in the target installation, an import may be performed, or other methods for merging the metadata may be used. In this regard, a procedure may be executed that initiates the merge step(s). Then, export-import functions may call sub-system import packages through standard routines that perform exports in their respective area.

[0082] Various packages are employed in the implementation of step 460. For example, changes to the transport set, such as modifying the import options, etc., are facilitated through functions in an application layer, which in turn perform all the database related operations through database layer. Various functions determine dependencies among objects to be exported and/or imported (due to foreign-key / master-detail / parent-child relationships).

[0083] A number of tables are involved in the implementation of the step 460. For example, a master table stores metadata about transport set. A master table also indicates whether object ACLs have been exported and need to be imported, indicates preferences for users/groups in ACL's, and indicates default options for import including whether to ignore warnings. A detail table stores information about explicit and implicit object selections and external references, and indicates the mode of import for explicit and implicit portal objects.

[0084] A log table stores information about what happened to each object being exported/imported, etc. as well as other pertinent data. Every system table that comprise the application being developed would have a corresponding mirror table called transport table, may be with some extra columns. These columns could typically store unique Ids to identify the transport sets and store the import mode of each object being imported.

[0085] In an embodiment, a mapping table stores the target mapping for a portal object. While merging data from the transport tables to the target tables in the target installation, a data resolution step may be implemented. Data resolution may be useful in a number of situations. For example, an object might not exist in the target installation. In this case, a new name (e.g., identity) is generated for the portal object, and an entry is made in the mapping table for the portal object to reflect the old name that was associated with the portal object in the source installation. If the object already exists in the target installation, then an entry is made to match the target installation's name to the source object in the source installation.

[0086] The information in the mapping table may be used to propagate the name changes to a portal object among related portal objects. For example, a change in the identity of a category of data items may be propagated to all the tables which reference the category (e.g., pages, items, etc.). In this regard, method 400 may include a step of propagation.

MIGRATION OF IMPLICITLY SELECTED PORTAL OBJECTS

[0087] According to one embodiment, mechanisms are provided for migrating implicitly selected portal objects during the migration process. For example, in an embodiment, a manifest screen is provided for the transport set. The manifest screen displays the explicitly selected object(s), implicit references, and external references. In an embodiment, a detail screen is provided for each explicitly selected portal object that displays the children, implicit references and external references for the explicitly selected portal object. Additionally, the external references may be moved to the explicit

objects list. What kind of objects fall under Explicit/Referenced/External are left to the end users to identify and implement accordingly.

[0088] According to one embodiment, a mechanism is provided to allow the mode of import to be set for the explicit and implicit objects in transport set. In an embodiment, the approaches allow the export of access control lists (ACL's) for the objects in the transport set. Also, the transport set provides a mechanism for the export of users and groups preferences when object ACLs are selected for export. In an embodiment, the approaches also provide a method of viewing the log of actions being performed.

Downloadable links to transport set -specific export-import scripts for Windows NT and UNIX may also be provided.

HARDWARE OVERVIEW

[0089] The approach for reconfiguring nodes described herein may be implemented in a variety of ways and the invention is not limited to any particular implementation.

The approach may be integrated into a computer system or a network device, or may be implemented as a stand-alone mechanism. Furthermore, the approach may be implemented in computer software, hardware, or a combination thereof.

[0090] FIG. 3 is a block diagram that depicts a computer system 300 upon which an embodiment of the invention may be implemented. Computer system 300 includes a bus 302 or other communication mechanism for communicating information, and a processor 304 coupled with bus 302 for processing information. Computer system 300 also includes a main memory 306, such as a random access memory (RAM) or other dynamic storage device, coupled to bus 302 for storing information and instructions to be executed

by processor 304. Main memory 306 also may be used for storing temporary variables or other intermediate information during execution of instructions to be executed by processor 304. Computer system 300 further includes a read only memory (ROM) 308 or other static storage device coupled to bus 302 for storing static information and instructions for processor 304. A storage device 310, such as a magnetic disk or optical disk, is provided and coupled to bus 302 for storing information and instructions.

[0091] Computer system 300 may be coupled via bus 302 to a display 312, such as a cathode ray tube (CRT), for displaying information to a computer user. An input device 314, including alphanumeric and other keys, is coupled to bus 302 for communicating information and command selections to processor 304. Another type of user input device is cursor control 316, such as a mouse, a trackball, or cursor direction keys for communicating direction information and command selections to processor 304 and for controlling cursor movement on display 312. This input device typically has two degrees of freedom in two axes, a first axis (e.g., x) and a second axis (e.g., y), that allows the device to specify positions in a plane.

[0092] The invention is related to the use of computer system 300 for implementing the techniques described herein. According to one embodiment of the invention, those techniques are performed by computer system 300 in response to processor 304 executing one or more sequences of one or more instructions contained in main memory 306. Such instructions may be read into main memory 306 from another computer-readable medium, such as storage device 310. Execution of the sequences of instructions contained in main memory 306 causes processor 304 to perform the process steps described herein. In alternative embodiments, hard-wired circuitry may be used in place

of or in combination with software instructions to implement the invention. Thus, embodiments of the invention are not limited to any specific combination of hardware circuitry and software.

[0093] The term “computer-readable medium” as used herein refers to any medium that participates in providing instructions to processor 304 for execution. Such a medium may take many forms, including but not limited to, non-volatile media, volatile media, and transmission media. Non-volatile media includes, for example, optical or magnetic disks, such as storage device 310. Volatile media includes dynamic memory, such as main memory 306. Transmission media includes coaxial cables, copper wire and fiber optics, including the wires that comprise bus 302. Transmission media can also take the form of acoustic or light waves, such as those generated during radio-wave and infra-red data communications.

[0094] Common forms of computer-readable media include, for example, a floppy disk, a flexible disk, hard disk, magnetic tape, or any other magnetic medium, a CD-ROM, any other optical medium, punchcards, papertape, any other physical medium with patterns of holes, a RAM, a PROM, and EPROM, a FLASH-EPROM, any other memory chip or cartridge, a carrier wave as described hereinafter, or any other medium from which a computer can read.

[0095] Various forms of computer readable media may be involved in carrying one or more sequences of one or more instructions to processor 304 for execution. For example, the instructions may initially be carried on a magnetic disk of a remote computer. The remote computer can load the instructions into its dynamic memory and send the instructions over a telephone line using a modem. A modem local to computer system

300 can receive the data on the telephone line and use an infrared transmitter to convert the data to an infrared signal. An infrared detector can receive the data carried in the infrared signal and appropriate circuitry can place the data on bus 302. Bus 302 carries the data to main memory 306, from which processor 304 retrieves and executes the instructions. The instructions received by main memory 306 may optionally be stored on storage device 310 either before or after execution by processor 304.

[0096] Computer system 300 also includes a communication interface 318 coupled to bus 302. Communication interface 318 provides a two-way data communication coupling to a network link 320 that is connected to a local network 322. For example, communication interface 318 may be an integrated services digital network (ISDN) card or a modem to provide a data communication connection to a corresponding type of telephone line. As another example, communication interface 318 may be a local area network (LAN) card to provide a data communication connection to a compatible LAN. Wireless links may also be implemented. In any such implementation, communication interface 318 sends and receives electrical, electromagnetic or optical signals that carry digital data streams representing various types of information.

[0097] Network link 320 typically provides data communication through one or more networks to other data devices. For example, network link 320 may provide a connection through local network 322 to a host computer 324 or to data equipment operated by an Internet Service Provider (ISP) 326. ISP 326 in turn provides data communication services through the worldwide packet data communication network now commonly referred to as the "Internet" 328. Local network 322 and Internet 328 both use electrical, electromagnetic or optical signals that carry digital data streams. The signals through the

various networks and the signals on network link 320 and through communication interface 318, which carry the digital data to and from computer system 300, are exemplary forms of carrier waves transporting the information.

[0098] Computer system 300 can send messages and receive data, including program code, through the network(s), network link 320 and communication interface 318. In the Internet example, a server 330 might transmit a requested code for an application program through Internet 328, ISP 326, local network 322 and communication interface 318.

[0099] The received code may be executed by processor 304 as it is received, and/or stored in storage device 310, or other non-volatile storage for later execution. In this manner, computer system 300 may obtain application code in the form of a carrier wave.

[0100] In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.